



Project Status Report

High End Computing Capability Strategic Capabilities Assets Program

October 10, 2018


Dr. Rupak Biswas – Project Manager
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

HECC Redefines Measure of Computer Resource Allocation and Usage



- The SBU Conversion Factor represents the difference in efficiency of a particular architecture in performing the workload on its Minimum Allocatable Unit (MAU), compared to the Pleiades baseline MAU.
- In May 2011, HECC provided allocation and usage information using a Standard Billing Unit (SBU) based on one node (12 cores) of Pleiades Westmere hardware.
- Under the Westmere SBU rate, Pleiades reached over 1.2 billion SBUs, and Electra—in less than two years of operation—reached more than 71.5 million SBUs.
- On October 1, 2018, a new SBU rate was released based on one node (28 cores) of Pleiades Broadwell hardware—the new rate better aligns SBUs to the type of hardware currently in production.
- The new SBU (SBU19) rate will be applied to past SBUs to provide a consistent view over history.

Mission Impact: The NASA Standard Billing Unit (SBU) is a way of standardizing work across dissimilar architectures. The new SBU definition helps users be more aware of the underlying computer architecture and true cost of computation—leading to more effective use of resources.



This graph shows the quarterly normalized usage in millions of SBU and allocation growth of HECC systems since the inception of the Westmere node Standard Billing Unit in 2011.

POC: Blaise Hartman, blaise.hartman@nasa.gov, (650)-604-2539, NASA Advanced Supercomputing Division

Applications Team Determines SBU Rate for NCCS Skylake Cluster



- The HECC Application Performance and Productivity (APP) team recently determined Standard Billing Unit (SBU) charging rates for the new Skylake-based nodes in Scalable Compute Unit 14 (SCU14) on the NCCS Discover cluster.
- Using six applications in the updated benchmark suite, the new rates reflect the relative performance on different node types compared to a baseline Broadwell node.
 - The previous SBU rates used a Westmere node as the baseline.
 - The HECC SBU suite comprises applications from space sciences (Enzo), CFD (FUN3D, OVERFLOW, USM3D), and climate/weather modeling (GEOS-5, nuWRF).
- As part of the rate change effective October 1 (see slide 2), APP staff calculated a factor of 0.254 to convert from old SBUs to new SBUs, which keeps the total HEC production capacity constant across the rate changes.

Mission Impact: By establishing charging rates that reflect the power of resources to compute the NASA workload, the agency is able to have a unified picture of high-end computing usage across different centers.



Scalable Compute Unit 14 (SCU14) on the Discover cluster at the NASA Center for Climate Simulation is a 20,800-core system capable of 1.6 petaflops, or 1.6 quadrillion floating-point calculations per second. After SCU14's integration, Discover will offer users a total of nearly 108,000 cores, with a combined peak performance of just over 5 petaflops. The HECC APP team determined an SBU rate of 1.61 for the Skylake-based nodes in SCU14, as they are 61% more productive than the baseline Broadwell node at the NASA Advanced Supercomputing facility.

POCs: Henry Jin, haoqiang.jin@nasa.gov, (650) 604-0165, NASA Advanced Supercomputing (NAS) Division;
Robert Hood, robert.hood@nasa.gov, (650) 604-0740, NAS Division, ASRC

Applications Team Begins Early-Access User Testing of Cloud Bursting



- The Application Performance and Productivity (APP) team began user testing of a cloud-bursting mechanism that seamlessly transfers jobs from Pleiades to run on resources in the cloud.
 - Users first annotate their job script with information about files that need to be staged to and from the cloud. They then submit the job from a Pleiades front-end node to an HECC Portable Batch System (PBS) server.
 - That server transfers the files to Amazon Web Services (AWS) and the job to a PBS server running in AWS. The AWS server then acquires cloud resources for running the job and starts it.
 - When the job is finished, the HECC PBS server copies the result files back from AWS.
- The early-access users selected for testing have jobs that need compute resources not currently available at the NAS facility; for example, nodes with NVIDIA Tesla V100 GPU accelerators. HECC allocated \$250,000 for cloud development and testing.
- Over the next nine months, the APP team will extend the cloud-bursting mechanism to include project allocations and accounting for resource consumption.

Mission Impact: The ability to seamlessly transfer jobs from HECC servers to the cloud ensures that NASA supercomputing services remain cost effective while providing access to resources that are not yet available on premises.



By transferring jobs from HECC front-end nodes to Amazon Web Services cloud resources, selected users now have an option for running on NVIDIA Tesla V100 GPU accelerators, greatly improving their science results.

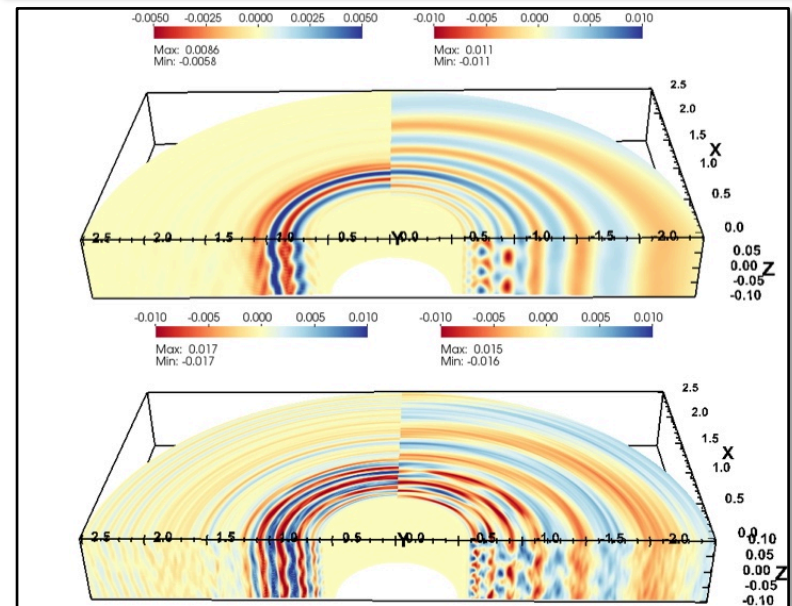
POCs: Steve Heistand, steve.heistand@nasa.gov, (650) 604-4369, and Sherry Chang, sherry.chang@nasa.gov, (650) 604-1272, NASA Advanced Supercomputing Division, ASRC

Access to Cloud-Based Resources Improves Modeling of Spiral Wave Instabilities



- Researchers studying spiral waves in protoplanetary disks were able to increase the resolution of their modeling four-fold through the use of GPU-accelerated nodes from Amazon Web Services (AWS).
 - A planet embedded within a gaseous disk around a young star excites spiral waves, which are subject to breaking into turbulence due to hydrodynamic instability.
 - In 2016, this spiral wave instability (SWI) was modeled at a resolution of $512 \times 128 \times 128$ —the limit due to large computational costs of the 3D simulation.
 - Using NVIDIA V100 GPUs on AWS, the increased resolution allowed the researchers to capture the cascade of SWI-driven turbulence at smaller scales.
 - The higher-resolution computation took 100 hours on a single AWS compute node with 8 V100s.
- The researchers are taking part in HECC's early-access user tests of cloud bursting, which seamlessly moves jobs from Pleiades to AWS (see slide 5).
- While the current simulations use a very simplified disk model, in the future the researchers plan to test the SWI case with more realistic models and higher resolutions.

Mission Impact: Access to NVIDIA Tesla V100 GPUs in the cloud allows researchers studying protoplanet dynamics to increase the resolution of their models without impacting their time to solution.



Top: Results with standard resolution ($512 \times 128 \times 128$). Bottom: Results with four times higher resolution ($2048 \times 512 \times 512$). While the spiral wave instability (SWI) develops with both numerical resolutions, the SWI-induced turbulence is better resolved with a larger number of grid cells.


POCs: Jaehan Bae, jbae@carnegiescience.edu, Carnegie Institute of Washington;
Gabriele Jost, gabrielejost@nasa.gov, (650) 604-0468, NASA Advanced Supercomputing Division, Supersmith

Machine Learning For Photovoltaic Cell Characterization



- The HECC Data Analytics team worked with the Photovoltaic and Electrochemical Systems Branch at Glenn Research Center to develop clustering and neural network models that help evaluate relationships in photovoltaic cell performance. Work includes:
 - Classified PV-Cell IV curves by chemistry.
 - Captured short-circuit current and open-circuit voltage.
- Over 31,000 IV curve data files were acquired and reduced down to a preliminary dataset containing ~6,500 IV curves and metadata. Findings include:
 - K-means clustering analysis of the IV curves suggests there are four types of IV curves, each with its own unique chemical signature.
 - A 2D convolutional neural network model trained on ~2800 samples of differing chemistry was able to calculate the short-circuit current and open-circuit voltage with 94% accuracy and a log mean square error of 0.02.
- This study helped correlate the electrical performance of solar cells with different chemical compositions. The next phase of the study will be to predict compositions for building high-efficiency solar panels.

Mission Impact: HECC is able to infuse machine learning technology into NASA research efforts to successfully demonstrate scientific results with faster time to discovery.



t-Distributed Stochastic Neighbor Embedding (t-SNE) of PV-Cell IV curves. Each curve has a unique chemical signature. (See key in notes.)

POCs: Shubha Ranjan, shubha.ranjan@nasa.gov, (650) 604-1918, NASA Advanced Supercomputing (NAS) Division;
Thaddeus Norman, thaddeus.j.norman@nasa.gov, (650) 604-2068, NAS Division, ASRC

Tools Team Upgrades HECC Incident Management System to BMC Remedy 18.05



- HECC's Tools team upgraded the Remedy incident, problem, and asset management system from version 9.1 to version 18.05.
- End user enhancements available with the Remedy 18.05 upgrade include:
 - License timeout and release, with user pre-notification, to provide better management of floating user licenses.
 - New features and functions in Smart Reporting.
 - Improved performance and stability when using hierarchical groups and full text searches.
- Remedy administration and performance enhancements include:
 - New Management Console to manage the Remedy server upgrades and patching.
 - Improved installers resulting in a 40% reduction in time-to-upgrade.
 - Effective management of settings under Centralized Configuration.
 - Reduced CPU usage for Remedy Mid Tier by 50%.

Mission Impact: The Remedy 18.05 upgrade provides HECC staff at the NASA Advanced Supercomputing facility with the latest Remedy features for administration tasks, and enhances the tool used to manage user calls and agency assets.

bmcsoftware AR System Deployment Management Console		
Manage Package	Filter By State <input type="radio"/> Draft <input type="radio"/> Deployable <input checked="" type="radio"/> All	
Create	4 of 4 results	
View	Package Name	Package Version
Delete	ITSM 9104 Patch	ITSM_9.1.04.002_180222
Copy	ITSM 9.1.04.002 Package for ALL platforms.	
	Atrium 9104 Patch	cmdb_9.1.04002_180221
	Package for ALL platforms.	
	ARServer 9.1.04 Patch	ARServer_9.1.04.002_20180416002
	Cumulative Hotfix for ALL platforms.	
	ARS 9104 Patch	ars_9.1.04002_20180217
	ARS 9.1.04.002 Package for ALL platforms.	
Operations		
Build		
Deploy		
Rollback		
Transfer Package		
Export		
Import		

The new Remedy Deployment Management Console allows for easy import, management, and rollback of Remedy upgrades and patches.

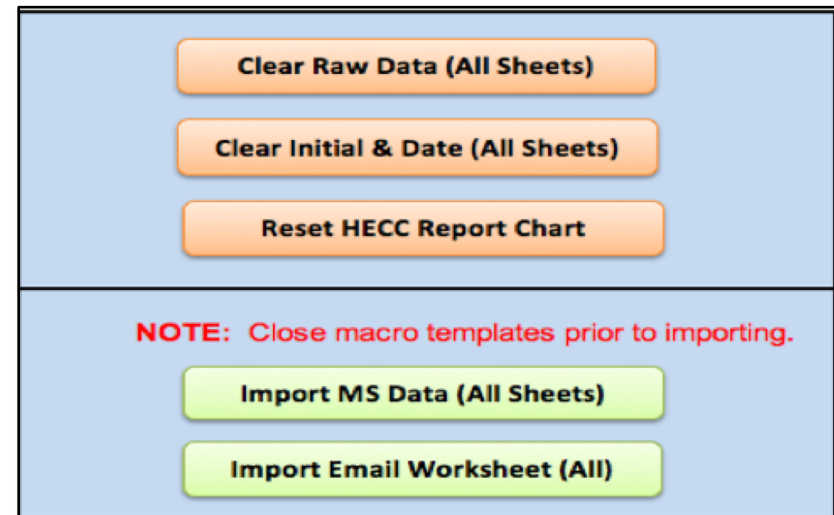
POC: Aqeel Rehman, mohammad.a.rehman@nasa.gov, (650) 604-4566, NASA Advanced Supercomputing Division, Intrinsix

Tools Team Continues Improvement of the HECC Monthly Reporting Process



- The Tools team continued to improve the monthly HECC reporting process by automating scripts, programs, and transfers of data to Excel spreadsheets.
- Recent improvements include:
 - Created Excel templates to efficiently format and import external data, reducing the need for manual input and removing numerous steps in the process; this condensed report development from 4–5 days to just 75 minutes.
 - Created script to summarize usage by program, project, and mission to verify data generated by a detailed GID report. This eliminated manual verification of the report.
 - Created script to generate CPU usage since 2008, using a static table containing information on usage for retired systems. Reporting time was reduced from 2 hours to 10 minutes.
- Tools staff continue to identify areas of improvement for both report generation and the addition of new systems into the reports.

Mission Impact: As the complexity of the HECC supercomputing environment continues to increase, automation of usage reporting helps to ensure accuracy and meet aggressive reporting schedules.



Using the reports template, the Tools team is able to import all HECC report data by clicking two buttons.

POCs: Mi Young Koo, mi.y.koo@nasa.gov, (650) 604-4528, and Sandra Glass, sandra.r.glass@nasa.gov, (650) 604-1416, NASA Advanced Supercomputing Division, ASRC

HECC Experts and Service Providers Meet with Users at NASA's Jet Propulsion Lab



- On September 19, a team of HECC supercomputing experts and service providers met with users at NASA Jet Propulsion Laboratory (JPL) as part of an outreach campaign aimed at improving the productivity of scientists and engineers supporting NASA missions.
- The meeting opened with a seminar, presented by HECC Deputy Project Manager William Thigpen, that was geared toward the broader HPC community. The seminar included an overview of the HECC project, accomplishments in advancing energy efficient computing, and a vision for the project's future.
- More than 40 people attended the seminar, including JPL HPC representatives. About 25 participated in user-focused presentations and discussion that followed. We found that:
 - Many users were not aware of all of the HECC services available to them.
 - Topics of interest included code optimization, visualization, cloud work, and how user success stories are selected and reported both to the public and up through NASA channels.
 - Users were also interested in the NAS Facility Expansion and Modular Supercomputing Facility, especially in regards to the water and power savings.
- A summary of these meetings, as well as findings and recommendations, will be provided in a report that is currently in progress.

Mission Impact: HECC systems and services support more than 200 NASA-sponsored scientists and engineers in Southern California. User outreach is imperative to ensure that their HPC requirements for completing their work are being met, and that these users are able to make the most of their computing time in support of NASA missions.

JPL SPACE | Upcoming Event

NASA blazes a different path to energy-efficient computing



JPL SPACE announcement for the September 19 seminar, which included information on the HECC project's revolutionary advancements in energy-efficient computing, as well as plans for the expansion of NASA's high-performance computing resources.

POC: Emily Kuhse, emily.kuhse@nasa.gov, (650) 604-1687, NASA Advanced Supercomputing Division, ASRC

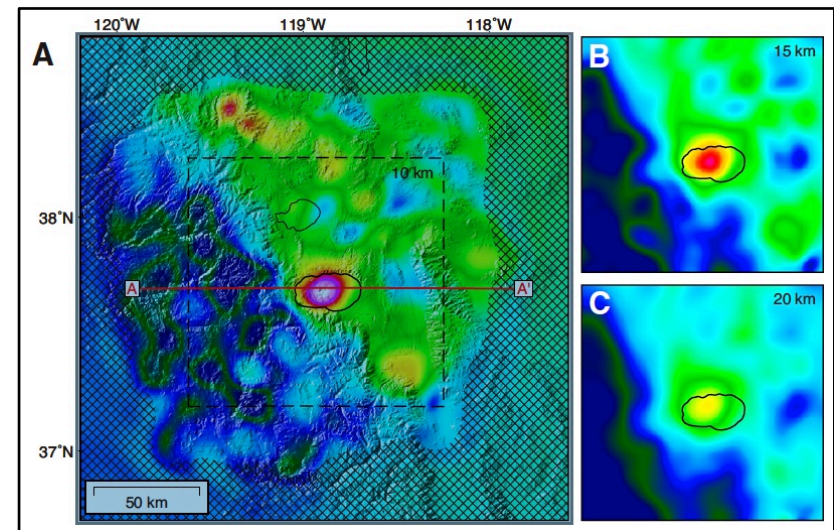
Pleiades Enables Discovery of Huge Magma Reservoir Beneath Supervolcano*



- A team of seismologists at the U.S. Geological Survey, University of California, Berkeley, and the University of Rhode Island used modeling and simulation methods on Pleiades to discover a 240-cubic-mile reservoir of magma under California's Long Valley Caldera.
- The researchers were able to image and map the region beneath the supervolcano with sufficient detail to determine the extent of the magma reservoir, which was not previously possible.
 - The team developed a method that solved for the caldera's crustal shear-wave velocity structure using 3D full-waveform tomography.
 - Source data included: Rayleigh-wave, ambient noise seismograms from all seismic stations within 150 km of the caldera over 26 years; and Rayleigh-wave records from 11 large regional earthquakes.
 - This capability allowed them to resolve depths of more than 20 km below sea level underneath the caldera, which was not possible with previous methods.
- The HECC Applications Performance and Productivity team worked closely with the researchers to optimize their code.

* HECC provided supercomputing resources and services in support of this work.

Mission Impact: Through access to HECC resources, this project furthers NASA's goal to better understand Earth's surface and interior to inform the assessment, mitigation, and forecasting of natural hazards such as earthquakes, tsunamis, landslides, and volcanic eruptions—one of the focus areas of the Science Mission Directorate.



Depth slices and west-east profile of the Long Valley tomographic model. A: Shear-wave velocity at 10 km below sea level (bsl). The caldera is outlined in thick black; Mono Lake is outlined in thin black. The dashed square shows the extents of two depth slices at 15 and 20 km bsl (B and C).

POC: Ashton Flinders, aflinders@usgs.gov, (650) 329-5050, United States Geological Survey, California Volcano Observatory

HECC Facility Hosts Several Visitors and Tours in September 2018



- HECC hosted 6 tour groups in September; guests learned about the agency-wide missions being supported by HECC assets, and some groups also viewed the D-Wave 2000Q quantum computer system. Visitors this month included:
 - Congresswoman Zoe Lofgren, U.S. Representative for California's 19th Congressional District, who toured the Modular Supercomputer Facility and the Quantum Artificial Intelligence Laboratory.
 - A group of senior executives from the Haas Business School at the University of California, Berkeley.
 - A group from the NASA Digital Transformation Formulation Team at Headquarters.
 - Håkan Bushke, President and CEO of aerospace and defense company SAAB Group, and Marcus Wallenberg, Chairman of the Board, SAAB; and Christer Fugelsang, Professor, Swedish Royal Institute of Technology, who met with NAS/HECC staff.



Research scientist Eleanor Rieffel (left) gives a technical review of the D-Wave 2000Q System to Congresswoman Zoe Lofgren of California's 19th congressional district.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputing Division



- **“Analysis of Transient Blow-Out Dynamics in a Swirl-Stabilized Combustor Using Large-Eddy Simulations,”** P. Ma, et al., Proceedings of the Combustion Institute, September 3, 2018. *
<https://www.sciencedirect.com/science/article/pii/S1540748918302499>
- **“Radiatively Induced Precipitation Formation in Diamond Dust,”** X. Zeng, Journal of Advances in Modeling Earth Systems, September 3, 2018. *
<https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001382>
- **“Dawnward Drifting Interchange Heads in the Earth’s Magnetotail,”** E. Panov, P. Pritchett, Geophysical Research Letters, September 4, 2018. *
<https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018GL078482>
- **“Superthermal Electron Acceleration in a Reconnecting Magnetotail: Large-Scale Kinetic Simulation,”** M. Zhou, et al., Journal of Geophysical Research: Space Physics, September 4, 2018. *
<https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JA025502>
- **“Metal Mixing and Ejection in Dwarf Galaxies is Dependent on Nucleosynthetic Source,”** A. Emerick, et al., arXiv:1809.01167 [astro-ph.GA], September 4, 2018. *
<https://arxiv.org/abs/1809.01167>

** HECC provided supercomputing resources and services in support of this work*

Papers (cont.)




- **“Modeling Photoionized Turbulent Material in the Circumgalactic Medium,”** E. Buie II, et al., The Astrophysical Journal, vol. 864, no. 2, September 5, 2018. *
<http://iopscience.iop.org/article/10.3847/1538-4357/aad8bd/meta>
- **“Assessment of Spray Combustion Models in Large-Eddy Simulations of a Polydispersed Acetone Spray,”** Q. Wang, et al., Proceedings of the Combustion Institute, September 5, 2018. *
<https://www.sciencedirect.com/science/article/pii/S1540748918301949>
- **“3D Simulations of Planet Trapping at Disc-Cavity Boundaries,”** M. Romanova, et al., arXiv:1809.04013 [astro-ph.EP], September 11, 2018. *
<https://arxiv.org/abs/1809.04013>
- **“Outbursts of Luminous Blue Variable Stars from Variations in the Helium Opacity,”** Y.-F. Jiang, et al., Nature: Letters, September 26, 2018. *
<https://www.nature.com/articles/s41586-018-0525-0>

** HECC provided supercomputing resources and services in support of this work*




- **NASA Leverages HPC for Asteroid Defense**, *HPCwire*, September 4, 2018—NASA has been actively working to understand and mitigate the risks posted by asteroids for a decade. Much of this work currently takes place in California as part of the Asteroid Threat Assessment Project (ATAP) at NASA Ames, using the Pleiades supercomputer.
<https://www.hpcwire.com/2018/09/04/nasa-leverages-hpc-for-asteroid-defense/>
- **NASA Administrator Meets Agency's Innovators in Silicon Valley**, *NASA Ames Press Release*, September 4, 2018—NASA Administrator Jim Bridenstine toured NASA's Ames Research Center for the first time on August 30, as the center prepares for its 80th anniversary next year. Bridenstine visited the NASA Advanced Supercomputing Division, which supports all NASA mission directorates.
<https://www.nasa.gov/feature/ames/nasa-administrator-meets-agencys-innovators-in-silicon-valley>
- **Video: Suitability of Commercial Clouds for NASA's HPC Applications**, *insideHPC*, September 10, 2018—Robert Hood from the NASA Advanced Supercomputing Division presented the results of an in-house study that asked if commercial cloud resources could be a cost-effective alternative to traditional supercomputing resources.
<https://insidehpc.com/2018/09/video-suitability-commercial-clouds-nasas-hpc-applications/>

HECC Utilization






September 2018

HECC Utilization Normalized to 30-Day Month




HECC Utilization Normalized to 30-Day Month




- ① 16 Broadwell racks added to Electra, 20 Westmere half racks added to Merope
- ② 4 Skylake E cells (16 D Rack Equivalence) added to Electra
- ③ 2 Skylake E cells (8 D Rack Equivalence) added to Electra; 1 rack is dedicated to ARMD

Tape Archive Status




September 2018


Tape Archive Status




Pleiades: SBUs Reported, Normalized to 30-Day Month




Pleiades: Devel Queue Utilization




Pleiades: Monthly Utilization by Job Length




Pleiades: Monthly Utilization by Size and Mission




Pleiades: Monthly Utilization by Size and Length




Pleiades: Average Time to Clear All Jobs




Pleiades: Average Expansion Factor




Electra: SBUs Reported, Normalized to 30-Day Month



Electra: Devel Queue Utilization




Electra: Monthly Utilization by Job Length




September 2018


Electra: Monthly Utilization by Size and Mission




Electra: Monthly Utilization by Size and Length




Electra: Average Time to Clear All Jobs



Electra: Average Expansion Factor




Merope: SBUs Reported, Normalized to 30-Day Month




*Utility Adjustment: Multiple failures of chillers in N233A necessitated turning off a large portion of Merope

Merope: Monthly Utilization by Job Length




September 2018


Merope: Monthly Utilization by Size and Mission




Merope: Monthly Utilization by Size and Length




Merope: Average Time to Clear All Jobs




Merope: Average Expansion Factor



Endeavour: SBUs Reported, Normalized to 30-Day Month




Endeavour: Monthly Utilization by Job Length




September 2018

Endeavour: Monthly Utilization by Size and Mission




September 2018

Endeavour: Monthly Utilization by Size and Length



Endeavour: Average Time to Clear All Jobs



Endeavour: Average Expansion Factor

